22.1 A Distributed Critical-Path Timing Monitor for a 65nm High-Performance Microprocessor

Alan Drake, Robert Senger, Harmander Deogun, Gary Carpenter, Soraya Ghiasi, Tuyet Nguyen, Norman James, Michael Floyd, Vikas Pokala

IBM, Austin TX

Scaling has caused an increase in process variation and the sensitivity of cycle time to workload and environmental conditions, making it difficult to predict the cycle time of microprocessors. Cycle time is determined by the required performance with an added timing margin determined by the acceptable yield. After manufacture, microprocessors are binned into performance categories to account for process variation, but because of the influence of the workload on cycle time, there is a danger of losing performance with overly conservative timing margins. A critical-path monitor (CPM) that measures critical-path delay and the effects of noise and localized V_{DD} droops on timing is designed as part of the POWER6™ microprocessor. The CPM also measures across-chip process variation and detects early life wearout mechanisms such as NBTI and hot-electron degradation.

Depending on the operating point, process corner, and workload of the microprocessor, many different timing paths may be critical. The CPM uses a small number of delay paths with different delay versus process, voltage, and temperature (DvPVT) curves to synthesize the critical paths. It is a time-to-digital converter that uses the system clock as the reference signal for conversion. The CPM, shown in Fig. 22.1.1, is composed of an edge-launching latch, delay-synthesis block, edge detector, data-analysis block, and configuration block. The POWER6 uses two CPM variations: one for core timing regions and one for nest (non-core) timing regions.

The synthesis block uses five delay paths for critical path synthesis: series 4-NAND gates, series 3-NOR gates, an adder path, a wire-dominated path, and series pass-gates. These paths are selected to have different DvPVT curves: the wire delay, adder, and pass-gate paths have different sensitivities to voltage, while the NAND and NOR paths have similar sensitivities to voltage as the adder path, but different sensitivities to process variation. Other timing monitors [1-3] vary the ratio of wire delay to MOS delay for their synthesis paths, but simulations indicate they may not capture the timing effect of tall PMOS and NMOS stacks. Simulations of delay lines with a mixed amount of wire delay show that delay lines with two- and three-wire segments have DvPVT curves equivalent to the wire-dominated path while delay lines with more than four wire segments have DvPVT curves equivalent to the adder path. The synthesis paths account for 60% of the cycle time in the core and 70% of the cycle time in the nest; the remaining time is used in the edge detector and switching circuitry.

There are two edge detectors in the CPM, based on [4] (Fig. 22.1.1). They consist of a 12-inverter delay line with capture latches connected to each inverter output. The minimum resolution of the edge detector is one FO2 delay. When an edge, *B*, enters the edge detector, its position is captured on the rising edge of the clock, ϕ. The latch output is then compared to the expected value of the edge detector, *A*, had the edge propagated all the way through the edge detector. The result is a thermometer code where the position of the 1-to-0 transition indicates how far the edge propagated into the edge detector. Timing slack is indicated by the difference between the current edge position and the calibrated edge position at failure.

The comparator provides for combining DvPVT curves. It selects the slowest path between the two edge detectors and the output register. This allows combining two different delay paths into one delay path; for example, the adder and wire paths can be selected so that the CPM follows the delay of the wire path at high voltages and the adder path at lower voltages.

The CPM operates as follows: on the rising edge of the system clock, an edge is launched into the synthesis paths; at the same time, the edge bypasses the synthesis block, resets the bits in the edge detector, and is inverted to provide the inverted edge for the next cycle; after passing through the synthesis path, the edge is latched in the edge detector on the rising edge of the system clock. The CPM captures rising and falling edge delay on alternating clock cycles. The core CPM is $90\times36\mu m^2$ and the nest CPM is $90\times48\mu m^2$ in 65nm SOI. There are 24 CPMs distributed across the microprocessor (Fig. 22.1.2): 8 in each core and 8 in the nest. Because of the time-to-digital nature of the output, its sensitivity to multiple variables, and its distribution across the microprocessor, the CPM measures local power-supply droop, clock instability, process variation, NBTI and early aging effects, and temperature in addition to timing; although, it is not always possible to separate these effects.

Figure 22.1.3 shows the measured delay versus voltage on nominal parts for the core CPM paths. The paths are normalized to demonstrate the different slopes of each delay path. There is some divergence in the pass-gate and wire delay paths from the MOS paths at the edges of the operating range. The small variation in the wire-delay path demonstrates that for even large percentages of wire delay, the MOS delay variation dominates the path delay.

Figure 22.1.4 shows the average maximum frequency of the microprocessor versus the measured bit position of the adder path for the CPMs in core 0 at each voltage. The curve is generated by running a heavy workload and increasing the frequency until failure. If the CPM exactly tracks the critical path, the bit position at failure should not change. There is an average of three bits of rise in the bit position as voltage rises, indicating the adder path does not exactly match the critical path. None of the paths exactly track the critical path, but because the output is a thermometer code, the difference is calibrated out by recording the bit position at failure, at two voltages and interpolating between them using a fitting function. Because wafer and module testing are done at different voltages, calibration adds little to the test time. Uncalibrated, the CPM tracks the critical path of the microprocessor to within 3 FO2 delays at the extreme voltages with a standard deviation of 1/2 a bit.

Figure 22.1.5 is a cycle-by-cycle graph of the CPM position in one core running a variable workload program, demonstrating the ability of the CPM to detect localized noise. The 9b spike occurs when the workload increases significantly. There is a 3b difference across the CPMs in response to the workload, indicating each CPM senses localized V_{DD} droop and other noise resulting from high switching activity.

To measure process variation, the CPMs are read with constant settings and no programs running. The CPMs are read in groups of 8 on the same clock cycle, so only clock skew and process variation affect the CPM timing. Figure 22.2.6 shows the delay variation across each of the cores and across the microprocessor for all delay paths at nominal voltage. The path delays exhibit a standard deviation of 2.3 to 4.8ps and have a maximum delay difference of 16ps. The CPM is measured to be sensitive to DC voltage changes as small as 10mV, and because it samples every clock cycle and can respond to any noise signal causing more than 1 FO2 change in path delay, it has a wide bandwidth.

Acknowledgements:
This work was supported in part by DARPA contract NBCH30390004. We gratefully acknowledge the contributions of J. Schaub, P. Restle, S. Taylor, and A. Saetow.

References:
[1] M. Elgebaly and M. Sachdev, "Efficient Adaptive Voltage Scaling System Through On- Chip Critical Path Emulation," *ISLPED*, pp. 375-380, Aug., 2004.
[2] S. Akui, K. Seno, N. Makai, et al., "Dynamic Voltage and Frequency Management for a Low-Power Embedded Microprocessor," *ISSCC Dig. Tech. Papers*, pp. 64-65, Feb. 2004.
[3] T. Fischer, J. Desai, B. Doyle, et al., "A 90-nm Variable Frequency Clock System for Power-Managed Itanium Architecture Processor," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 218-228, Jan., 2006.
[4] P.J. Restle, R.L. Franch, N.K. James, et al., "Timing Uncertainty Measurements on the POWER5 Microprocessor," *ISSCC Dig. Tech. Papers*, pp. 354-355, Feb., 2004.

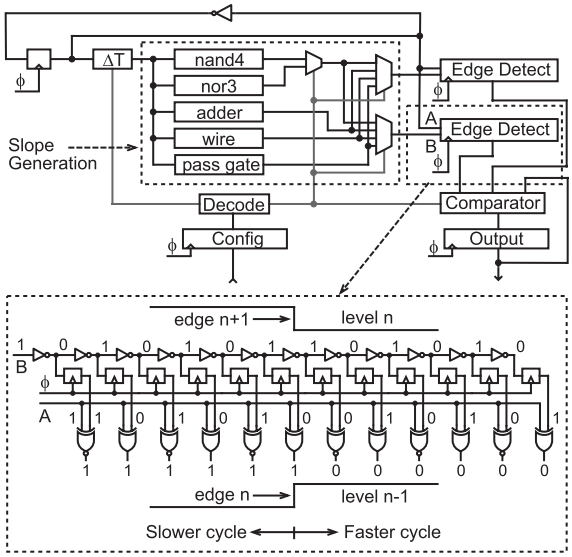

Figure 22.1.1: Block diagram of the CPM and edge detector.

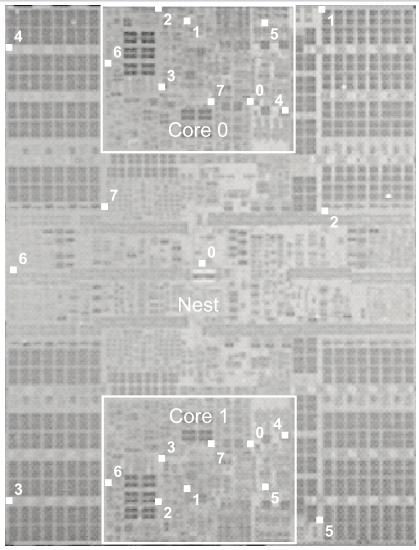

Figure 22.1.2: Die shot of the POWER6 microprocessor showing the location of the CPMs.

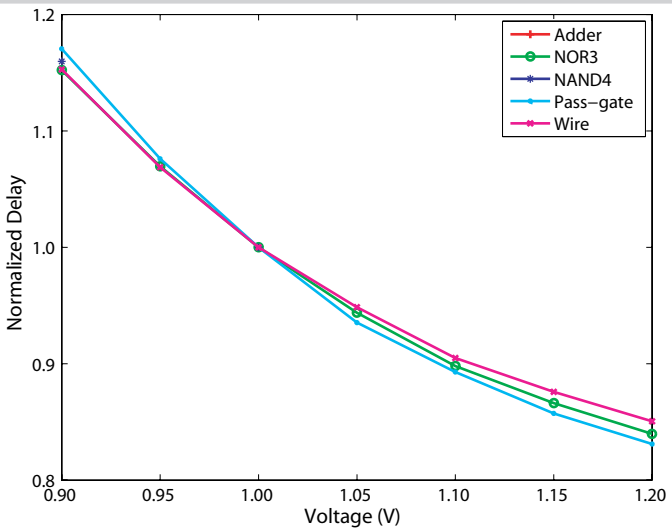

Figure 22.1.3: Normalized delay, from edge launch to the center bit position of the edge detector, of CPM paths in core 0 in nominal hardware.

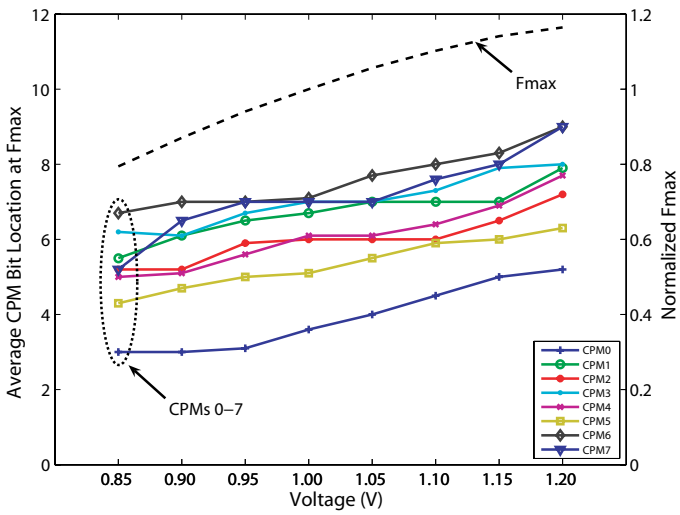

Figure 22.1.4: Core 0 adder path delay at the maximum processor frequency.

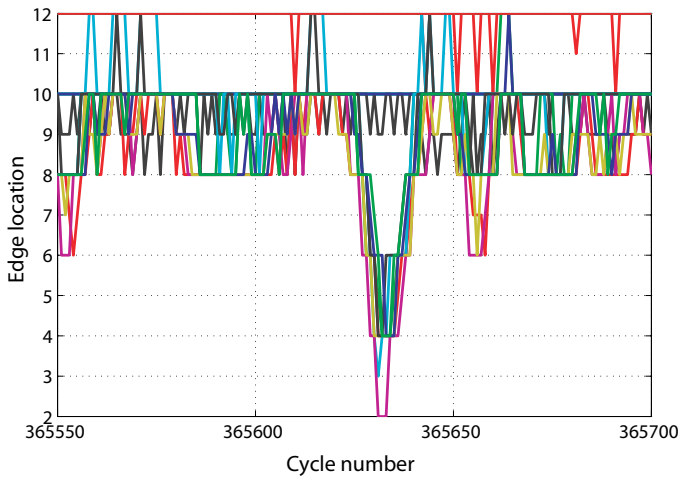

Figure 22.1.5: Cycle-by-cycle variation in CPM bit position of the adder path for all 8 CPMs in core 1 as a function of work-load.

	adder	nor	nand	pass-gate	wire
core 0: σ	3.78	3.45	3.91	3.28	2.82
core 1: σ	3.55	3.16	2.25	3.55	4.21
chip: σ	4.09	4.03	3.90	4.80	4.10
core 0: Δ	10	9	11	10	7
core 0: Δ	12	9	6	10	12
core 0: Δ	13	11	14	16	15

Figure 22.1.6: Variation, in ps, of the CPM path delay measured in each core and across both cores at nominal voltage.